

Tools for verifying and assessing the validity of social media and user-generated content



(veri.ly)

By [Josh Stearns](#) and [Leighton Walter Kille](#)

This post was adapted from [VerificationJunkie](#), a site created by [Josh Stearns](#) of the Geraldine R. Dodge Foundation. [@jcstearns](#)

“Interesting if true” is the old line about some tidbit of unverified news. Recast as [“Whoa, if true”](#) for the Twitter age, it allows people to pass on rumors without having to perform even the most basic fact-checking — the equivalent of a whisper over a quick lunch. Working journalists don’t have such luxuries, however, even with the continuous deadlines of a much larger and more competitive media landscape. A cautionary tale was the [February 2015 report](#) of the death of billionaire Martin Bouygues,

head of a French media conglomerate. The news was instantly echoed across the Web, only to be [swiftly retracted](#): The mayor of the village next to Bouygues's hometown said that "Martin" had died. Alas, it was the wrong one.

The issue has become even knottier in the era of [collaborative journalism](#), when nonprofessional reporting and images can be included in mainstream coverage. The information can be crucial — but it also can be wrong, and even intentionally faked. For example, two European publications, *Bild* and *Paris Match*, said they had seen a video purportedly shot within the Germanwings flight that crashed in March 2015, but doubts about such a video's authenticity [have grown](#). (Of course, there is a [long history of image tampering](#), and news organizations have been culpable year after year of running — and even producing — manipulated images.)

The speed of social media and the sheer volume of user-generated content (UGC) make fact-checking by reporters even more important now. Thankfully, a wide variety of digital tools have been developed to help journalists check facts quickly. This post was adapted from [VerificationJunkie](#), a directory of tools for assessing the validity of social-media and user-generated content. The author is [Josh Stearns](#), director of the journalism sustainability project at the Geraldine R. Dodge Foundation.

Research, case studies

[The Verification Handbook](#): A guide to verifying digital content for emergency coverage authored by journalists from the BBC, Storyful, ABC, Digital First Media and others. Released under the Creative Commons license, it provides tools, techniques and step-by-step guidelines for how to deal with user-generated content during emergencies. PBS.org review: ["Verification Handbook Mixes Tools, Tips and Culture for Fact-Checking."](#)

[The Verification Handbook for Investigative Reporting](#): A follow-up to *The Verification Handbook*, this guide highlights techniques for leveraging user-generated content and open-source information in investigative reporting. Subjects include how to use databases, domain records and materials to investigate companies; verifying data quality; building expertise through UGC verification; and applying ethical principles to digital investigations. There are also three case studies.

BBC Verification Hub: "Started in 2005 to sift through unsolicited contributions previously perused by many different teams, the [BBC Verification Hub] has grown to a complement of 20 staffers," writes David Turner in a [NiemanLab article](#). "Initially the team focused heavily on images, footage and eyewitness accounts e-mailed to the BBC, but in the past few years people have become much more prone to distribute material themselves through Twitter, YouTube and Facebook. As a result, the number of contributions proffered to the BBC has declined to about 3,000 a day, and the Hub's task has moved toward semi-conventional newsgathering with a Web 2.0 twist. Staffers now use search terms, see what's trending on Twitter, and look at the images and footage trusted contacts are discussing on their Twitter streams."

“Rumor Cascades”: Data scientists from Facebook and Stanford University prepared this paper for the 2014 International Conference on Weblogs and Social Media. The researchers, Adrien Friggeri, Lada A. Adamic, Dean Eckles and Justin Cheng, identify known rumors through Snopes.com (the urban legend reference site) and analyze how nearly 4,800 distinct rumors circulated on Facebook. Among those studied, 22% were related to politics and 12% involved fake or doctored images. False rumors thrive on Facebook: In the Snopes database, 45% of rumors are false, while 26% are true; in contrast, on Facebook 62% of rumors are false and only 9% are true. The authors note that “true rumors are more viral — in the sense that they result in larger cascades — achieving on average 163 shares per upload whereas false rumors only have an average of 108 shares per upload.” Even when people discover the falsity of a rumor and delete their reshare, it does not appear to affect the unfolding cascade. The “popularity of rumors — even ones that have been circulating for years in various media such as email and online social networks — tends to be highly bursty. A rumor will lie dormant for weeks or months, and then either spontaneously or through an external jolt will become popular again.”

“How to Separate Fact and Fiction Online”: 2012 TED talk from Markham Nolan of Storyful. In this case study, a Storyful team verifies a user-generated YouTube video of lightning hitting a tree using only free Web tools. Starting with the name of the user who uploaded the video, the team verifies the information on Spokeo, cross-references those results with a weather report via Wolfram Alpha, tracks down an exact address in the White Pages, and uses Google Maps satellite images to match the house and yard in the video to the address. Storyful article: [“The Human Algorithm.”](#)

“Location-based Trust for Mobile User-generated Content”: In this 2008 study, researchers from the Department of Electrical Engineering at Princeton University explore how to establish the authenticity of content created by “untrusted mobile users.” Using “secure localization and certification service,” they develop and propose a tool that would help content producers tag their content with a spatial timestamp indicating its physical location while also protecting privacy.

“Automatically Identify Fake Images on Twitter”: In a 2013 paper from the Indraprastha Institute of Information Technology, IBM Research Labs and the University of Maryland, the researchers found that it was possible to identify tweets containing fake Hurricane Sandy images with up to 97% accuracy. The paper provides interesting data about the way fake images spread during Sandy, and explores how one day we may be able to flag tweets as potentially containing false information. Poynter.org article: [“New Research Suggests It’s Possible to Automatically Identify Fake Images on Twitter.”](#)

Dynamic Network Analysis: On Jan. 25, 2011, the day that Egyptian president Mubarak was forced out of office, André Panisson created a real-time infographic mapping tweets and retweets. “As a tool for verification it helps you see the flow of information, or misinformation and track it back to its source. In addition, it helps you access who influential people are in a discussion, offering you leads and potential sources.” Panisson described the project this way: “It was very interesting to see, in real time, the exact moment when Tahrir Square, from a mass protest demonstration, has been transformed in a giant party, and the burst in the Twitter’s activity. It was like covering in real time a virtual event, a big event that was happening in the Twitter virtual world.” Article in *Fast Company*: [“Infographic of the Day: Watch Egypt’s Twitter Uprising Bloom.”](#)

Fact-checking sites

[FactCheck.org](#): A project of the Annenberg Public Policy Center of the University of Pennsylvania, the site is a “nonpartisan, nonprofit ‘consumer advocate’ for voters that aims to reduce the level of deception and confusion in U.S. politics.” While its focus is on politics, that topic is taken broadly and encompasses a lot of Web content.

[Checkdesk](#): A verification tool designed to help curate user-generated content during breaking news and connect journalists to citizen sources on the ground. “Checkdesk facilitates collaborative fact-checking of unverified reports,” the developers write. “Professional journalists can join forces with citizen journalists in search of background information and evidence to corroborate social media reports.” Introduction from Meedan.org: [“Checkdesk: A New Approach to Fact-checking Citizen Media of the Arab Spring.”](#)

[Full Fact Finder](#): This U.K.-based site covers information on the economy, health, crime and the law, immigration and education. “Search results offer users general background information, as well as details on the sort of data available in the area and links to statistics from official bodies.” Coverage from journalism.co.uk: [“Full Fact launches Online Fact-finding Tool.”](#)

[Emergent.Info](#): The site’s tagline is “real-time rumor tracker.” For example, on April 1, 2015, it checked whether a man was wanted in England for slapping people who sneezed in public (true) and a claim that doctors had confirmed the first death due to genetically modified food (false). The site is part of a research project of the Tow Center for Digital Journalism at Columbia University that focuses on how unverified information and rumor are reported in the media. Article from Craig Silverman, a fellow at the Tow Center: [“Researching Rumors and Debunking for the Tow Center at Columbia University.”](#)

[Churnalism](#): From the Sunlight Foundation, Churnalism is based on a U.K. site and compares articles to a database of press releases. It’s intended as a public-accountability tool but could also be useful for journalists assessing blog posts and other source material. Poynter article: [“Sunlight Foundation’s New Plagiarism-detection Software Launches, Claims a Bust.”](#)

[LazyTruth](#): An inbox extension that recognizes emails full of political myths, urban legends or security threats and debunks them in your mail program. It currently works only in Chrome and Gmail, but may be expanded to other browsers. Lifehacker article: [“LazyTruth Fact Checks Chain Emails, Responds to the Sender with the Truth.”](#)

Is Twitter Wrong: A [Tumblr](#) and [Twitter](#) account dedicated to debunking misinformation on Twitter.

Fact-correction tools

MediaBugs: A service for reporting specific, correctable errors and problems in media coverage. “We’ll provide a neutral, civil, moderated discussion space,” they state. “We’ll try to alert the journalists or news organization involved about your report and bring them into a conversation. As a result of this dialogue between journalists and the public, some errors may get corrected; others won’t. Either way, the discussion will leave a useful public record.” NiemanLab article: [“MediaBugs Rethinks Corrections by Taking a Page from Programmers.”](#)

Retwact: A tool that automates the process of notifying anyone who retweeted an inaccurate tweet from your account; the goal is to help slow the spread of misinformation by making it easier to correct tweets. *Atlantic* article: [“Retwact: A Tool for Fixing Twitter’s Misinformation Problem.”](#)

Report an Error Alliance: This is an ad-hoc group of individuals and organizations who endorse the idea that websites should always have an easy-to-find and -use “report an error” button. It’s a way of saying to users that you care about accuracy, you want to know when you make errors, and you’re conscientious about fixing them.

Image tools

TinEye: A reverse-image search engine, TinEye allows you to find out where an image came from, how it is being used, if modified versions exist, or to find higher-resolution versions. TinEye is the first image search engine to use image-identification technology rather than keywords, metadata or watermarks. It is free to use for non-commercial searching. IJNET article: [“Journalist’s Guide to Verifying Images.”](#)

Google Images: With Google Images’ “Search by Image” option you can upload an image and Google will show you any images that resemble it. It is a quick way to easily track down original source images, or spot modifications and edits to an image.

FourAndSix: An extension for Adobe Photoshop, FourMatch analyzes open JPEG images to determine whether they are untouched originals from a digital camera. As of April 2015 the service had been discontinued, however, and one called [Izitru](#) was recommended. FourAndSix article on fake photos, many of which were run by media organizations: [“Photo Tampering through History.”](#)

Reporting tools

SRSR: This application — its name stands for “Seriously Rapid Source Review” and is pronounced “sourcer” — helps aggregate and assess sources on social media during breaking news events. “The team built in custom computations and cues designed to assess potential sources based on location, network and past content. The app helps with identifying eyewitnesses and user archetypes, and visually cueing location, network and entities. Introductory blog post: [“Finding News Sources in Social Media.”](#) Academic paper: [“Unfolding the Event Landscape on Twitter: Classification and Exploration of User Categories.”](#)

SwiftRiver: A platform that helps people make sense of a lot of information in a short amount of time, SwiftRiver enables the filtering and verification of real-time data from channels like Twitter, SMS, Email and RSS feeds. The SwiftRiver platform offers organizations an easy way to apply semantic analysis and verification algorithms to different sources of information. SwiftRiver offers an open source, affordable, data intelligence platform for news organizations, non-profits, small governments, and NGOs. Coverage from Harvard’s Nieman Lab: [“‘Adding context to content’: Swift River gets Knight funding to tackle the problem of real-time verification.”](#)

WhoWhatWhen: A database of people and events from 1000 A.D. to the present that can be sorted, compared and aligned to confirm the accuracy of references to time, people and events. You can create graphic timelines that provide context for events and people’s lives. Good for confirming if a technology or world event actually happened during someone’s lifetime.

MemeTracker: Builds maps of the daily news cycle by analyzing around 900,000 news stories and blog posts per day from 1 million online sources, ranging from mass media to personal blogs. The site tracks the quotes and phrases that appear most frequently over time and thus can be useful for tracing the spread of misinformation. The program can help show how certain stories persist while others fade quickly. Ethan Zuckerman blog post: [“Jure Leskovec on Memetracker, Quantitative Media Analysis.”](#)

Keepr: Still in development, this social-media monitoring tool is intended to help journalists stay up to date with breaking news on Twitter. Developed around linguistic analysis, the tool helps mine information and trends in real time, highlighting relationships between people and information.

Citizen Desk: Still in development, this Web platform allows citizens and mobile journalists to send reports via SMS. “These enter the Citizen Desk, where they are verified on the spot by editors and published to a live blog. Editors are able to easily search YouTube, Twitter, Flickr and Google from within the app and then add these sources as supporting material, alongside their own text or images. All this is displayed on a light webpage which adapts to Web, tablet and mobile browsers. Users can comment on the live-blog using Facebook comments to discuss veracity and providence.” Emergency Journalism case study: [“Helping Citizen Journalists Cover Mozambique’s Elections.”](#)

InformaCam: A mobile app from the Guardian Project, it allows Android devices to embed images and videos with geotemporal and other metadata that will help others verify their authenticity. Users can then sign them with a digital signature unique to the device's camera sensor, encrypt and then send those files to someone they trust who maintains a secure server. Among other elements, the process preserves the chain of custody of the media, making it more likely to be admissible in a court of law. Walkthrough: ["Is This For Real? How InformaCam Improves Verification of Mobile Media Files."](#)

Ifussss: The name of this site stands for "If You See Something Share Something" focuses on monetizing verified, user-uploaded content. [10,000 Words](#) reports "You see traffic on a bridge, for example. You shoot and upload it to the ifussss network. It's automatically geo, time and hash-tagged. News editors can search and monitor the ifussss newsroom platform and, this is where it gets interesting, buy the content."

Veri.ly: A platform to crowdsource verification and crisis information, the service is not yet publicly available. The system is rooted in research by DARPA on distributed networks for finding specific information. *MIT Technology Review* article: ["Preventing Misinformation from Spreading through Social Media."](#)



A project of the Harvard Kennedy School's [Shorenstein Center](#) and the [Carnegie-Knight Initiative](#), Journalist's Resource is an open-access site that curates scholarly studies and reports.